

DE LA RECHERCHE À L'INDUSTRIE



www.cea.fr

Keynote Address


Storage for Exascale: To Infinity and Beyond

Philippe DENIEL (philippe.deniel@cea.fr)



Welcome to Paris!

Paris will host a prestigious event in 2024...



...but Paris currently hosts (and has hosted) even more prestigious events since 2011



(No logo for 2011)

Data Storage: an archeological approach

Data Storage for HPC has evolved a lot since the 90's

■ 90's: The HPC Cavemen's Era

- Cray T924 and Cray T3E, $\sim 60\text{Gflops}$ ($6 \cdot 10^7$ flops), 80 TB tapes
- Fun fact: Today's Samsung Galaxy S10+ = 8×13.4 Gflops




■ 2000: Terascale period (beginning of the SMP Cluster Age)

- Tera1 supercomputer installed at CEA
- 5 Tflops ($5 \cdot 10^9$ flops)
- 50TB (disks), 1PB (tapes)
- Fun fact: STK 9840 "Eagle" tapes had 40GB capacity, we had 25000 of them...



■ 2005: First Intermediate Period

- Tera10 supercomputer @ CEA
- 60 Tflops ($6 \cdot 10^{10}$ flops), 2 PB (disks, 100 GB/s), 10 PB (tapes)



■ 2010: Petascale period (beginning of the "Cluster of clusters" Age)

- Tera100 supercomputer @ CEA
- 1.05 Pflops (10^{12} flops), 20 PB (disks, 500 GB/s), 30 PB (tapes)



■ 2015: Second Intermediate Period

- Tera1000 supercomputer @ CEA, currently in production
- 30 Pflops ($3 \cdot 10^{13}$ flops)
- 40 PB (HDD, 767 GB/s), 2 PB (SSD, 1 TB/s), 150 PB (tapes)



Today... and Tomorrow

~~Winter~~ Exascale is coming

- 2020: Early exascale supercomputer
 - EXA1 to be installed at CEA in 2020-2021
 - Hundreds of Pflops to 1 Eflops (10^{14} flops to 10^{15} flops)
 - RFP for storage systems just began
 - No fun fact: I can't tell more about it or I will have to kill you afterward...



Expected requirements for Exascale

- In 2024, there will be supercomputers with up to 1-10 Eflops, with “warp speed” requirements...
 - 20-30 PB of RAM
 - Tens of millions of compute cores
 - 200 000 – 300 000 nodes: each of them are FS clients!!!
 - From 20 EB ($2 \cdot 10^{19}$ bytes) to 200 EB ($2 \cdot 10^{20}$ bytes) disks
 - From 100 EB to 1 ZB (10^{21} bytes) tapes
 - Fun fact: internet's storage size was estimated 32 ZB in 2018...
- Three Technological Challenges:
 - The memory/core ratio is strongly reducing because of “manycores” CPU
 - The Volume of Data and Metadata to be managed becomes really huge
 - Get prepared to manage several hundreds of thousands of clients...



There will be exciting challenges for Lustre!

Thank you again for coming to paris

I hope that you will enjoy LAD'19

- I am sure it will be even better than the 8 earlier editions!
- Many thanks to Florent, Gaël, Guillaume, Jérôme, Karine, Laetitia, Quentin & Thomas who made this event possible



Special thought for attendees who still suffer from jetlag

- Caffeine is your friend!

